High-level Reasoning and Low-level Learning for Grasping: A Probabilistic Logic Pipeline

Laura Antanas, Plinio Moreno, Marion Neumann, Rui Pimentel de Figueiredo, Kristian Kersting, José Santos-Victor, Luc De Raedt

arXiv:1411.1108v1 [cs.RO] 4 Nov 2014

***Abstract*—While grasps must satisfy the grasping stability criteria, good grasps depend on the specific manipulation scenario: the object, its properties and functionalities, as well as the task and grasp constraints. In this paper, we consider such information for robot grasping by *leveraging manifolds* and *symbolic object parts*. Specifically, we introduce a new *probabilistic logic* module to first *semantically reason* about pre-grasp configurations with respect to the intended tasks. Further, a mapping is learned from part-related visual features to good grasping points. The probabilistic logic module makes use of object-task affordances and object/task ontologies to encode rules that generalize over similar object parts and object/task categories. The use of probabilistic logic for task-dependent grasping contrasts with current approaches that usually learn direct mappings from visual perceptions to task-dependent grasping points. We show the benefits of the full probabilistic logic pipeline experimentally and on a real robot.**

***Index Terms*—task-dependent robot grasping, probabilistic logic pipeline, semantic grasping, local shape grasping**

I. Introduction

While robot vision capabilities are essential for perceiving and interpreting the world, robot grasping skills are essential for acting in arbitrary and dynamic environments and executing object manipulation tasks. Objects can be grasped in different ways. While performing a grasp, we must, at least, satisfy the grasping stability criteria, performing a good grasp also depends on the specific manipulation scenario: the object, its properties and functionalities, as well as task constraints and grasp constraints (e.g., gripper configuration). How to take into account such information for grasping is exactly the question we tackle in the present paper.

Specifically, instead of just learning a function that directly maps visual perceptions to task-dependent grasps, we introduce, as key contribution, an intermediary *probabilistic logic* module to *semantically reason* about the most likely object part to be grasped, given the object properties and task constraints. Then, a mapping is learned from part-related local visual features to good grasping points. The introduced symbolic part-based representation has several advantages:

- grasp transfer to novel objects that share similar parts and thus, generalization over similar (multiple) object parts;
- high-level task-dependent reasoning over parts reduces possible grasps and hence, improves performance;
- the use of symbolic parts as manifold information for reliable object category estimation.

The present paper investigates robot grasping by leveraging symbolic world knowledge, in the form of object/task ontologies and object-task affordances, visual features, object categorical and task-based information, by integrating them into one *probabilistic logic pipeline*. World knowledge and relations are encoded in compact logical grasping models that generalize over similar object and task categories, thus the logic offers a natural way to encode high-level knowledge. However, often, descriptions of the perceived world are also uncertain. For example, not all cups look like the 'prototypical' cup. Thus, we need probabilistic models, which, additionally, allow one to reason about the uncertainty in the world. We assume we have observations about the task and visual scene available and we show that, using our probabilistic logic, we can ask queries about different grasping aspects.

This paper makes the following contributions:

1) the integration of object category and task information for semantic object part grasping,
2) a first probabilistic logic module for task-dependent reasoning about robot grasping,
3) a general rule-based model encoding object-task affordances and objects/tasks ontologies that reflect world knowledge and allow generalization over object/task categories.

By employing object-task affordances and objects/tasks ontologies, the proposed pipeline can generalize over similar object parts and object/task categories. This allows us to experiment with a wide range of object and task categories, which is a critical aspect of autonomous agents acting freely in new environments. Our approach can be extended beyond the set of categories used, by augmenting the probabilistic logic module with extra rules. The benefit of introducing the proposed probabilistic logic pipeline is shown experimentally.

A. The proposed pipeline

The pipeline is exemplified in Fig. 1. It takes as input the point cloud of an object (i.e., a cup) and, using vision-based techniques, we first obtain a description of the scene in terms of symbolic object parts, object pose and containment. We assess global object similarity via manifold-based graph kernels to complete the scene description with a prior on the object category. Next, using the visual description, we query the probabilistic logic module for the most likely object category, most likely task and best pre-grasp. For our cup example, the manifold shape model predicts the categories *cup*, *can* and *pot* with probabilities 0.56, 0.36 and 0.05, respectively. The categorical logic module reasons about the symbolic parts and recalculates the probabilities as following:

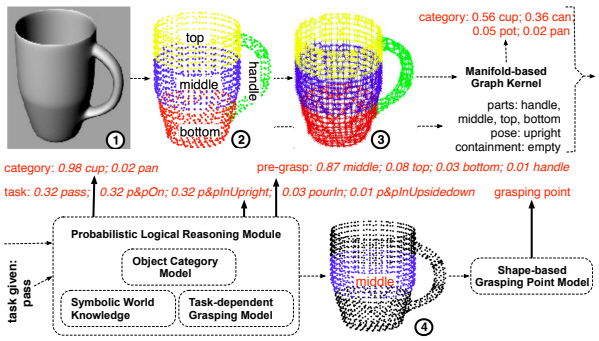

Fig. 1: The task-dependent grasping pipeline on a *cup* point cloud example. Top row (left to right): object ①, symbolic object parts ② with labels *top* (yellow), *middle* (blue), *bottom* (red), and *handle* (green), k-nn graph ③ with part labels, $k = 4$ (the edges are colored according to the colors of the adjacent nodes), manifolds model with its outcome and visual description of the object (pose, containment and parts). Bottom row: probabilistic logic module with its components and reasoning outcome, predicted pre-grasp *middle* ④, shape-based grasping model and predicted grasping point.

0.98 for category *cup* and 0.02 for category *pan*. The presence of exactly one *handle* increases the probability of the object being more a *cup* rather than a *can* and identifies the object more as a *pan* rather than a *pot*. Similarly, using object-task affordance knowledge and object/task ontologies (e.g., any object affords the tasks *pass* and *pick-place on*), but also world knowledge (e.g., the task *pour in* cannot be executed on a full object), the probabilistic logic module predicts the tasks *pass*, *pick-place on* and *pick-place inside upright* with equal probability. If the task given is *pass*, the task-dependent grasping model predicts as most likely pre-grasp the *middle* part of the object. The last step in the pipeline is the local shape-based grasping module which predicts the best point for grasping in the pre-grasp point cloud.

An overview of our learning and reasoning grasping pipeline is shown in Fig. 2. It has four modules. The first module is a visual perception module (detailed in Section III) which maps the object point cloud to scene descriptions in the form of symbolic and probabilistic visual observations about the world. After it segments the object point cloud and performs a full object shape reconstruction, the visual module estimates the object pose (*upright*, *sideways*) and parts (*top*, *middle*, *bottom*, *handle*, and *usable area*). Further, it predicts a prior on the object category by employing object similarity based on manifold and semantic part information via a graph kernel. The second module is responsible for pre-grasp prediction. It consists of a probabilistic logic reasoning model for task-dependent grasping which, given the input observations, is able to perform inference about the grasping scenario. The module can answer, in turn, queries about the object category, most likely task and most likely object part to grasp, given the task. It uses probabilistic visual observations about the object, such as object pose, object category (i.e., a prior on the object category predicted by the aforementioned graph kernel-based probabilistic estimator), and object functionality (i.e., the object is empty or full) provided by the first module, and evidence about the task, either given by a user, or in the form of a probabilistic prior on the task type. A more detailed explanation of this module is provided in Section IV. Once we have identified the most likely object part to grasp, the pipeline calls the third module, which solves the problem of grasping point prediction using local shape features of the object part. The description of this module follows in Section V. Finally, the last module takes the best predicted grasp and executes it on the robotic platform.

We proceed as follows. We start by reviewing the related work. Next, we explain in more detail each component of the proposed pipeline: the vision-based module, the probabilistic logic module and the grasping point prediction. Before concluding, we present our experiments in simulation as well as on a real robot.

II. Related Work

A. Visual-dependent grasping

The majority of grasping methods consider mainly visual features to learn a mapping from 2D/3D features to grasping parameters [1], [2], [3], [4]. Nevertheless, these methods have a major shortcoming: it is difficult to link a 3D gripper orientation to solely local image features. Only recently, methods that take global information into account have been proposed [5], [6]. The benefit is an increased geometric robustness, which is advantageous with respect to the pre-shape of the robotic hand and the general shape of the object, generating more accurate grasps. However, global information relies on a complete shape of the object.

In recent years shape completion using a single view has been extensively studied, typically in robotics grasping applications. Usually multiple object partial views are acquired from different viewpoints, using 3D range cameras, and the gathered point clouds are then registered and aligned together in a common reference frame. The Iterative Closest Point algorithm [7] and efficient variants [8] are often used to compute the alignment transformations and to build a complete object shape model [9]. However, when only a single view is available and/or it is not possible to acquire several views due to time constraints or scenario/robot restrictions the shape completion problem becomes harder and some assumptions

or pattern analysis must be made. In this direction, a wide range of ideas have been proposed including fitting the visible object surface with primitive shapes such as cylinders, cones, parallelepipeds [10], [11] or with more complex parametric representations like superquadrics [12].

Closely related to our shape completion approach, Thrun and Wegbreit [13] proposed a method based on the symmetry assumption. This method considers 5 basic and 3 composite types of symmetries that are organized in an efficient entailment hierarchy. It uses a probabilistic model to evaluate and decide which are the completed shapes, generated by a set of hypothesized symmetries, that best fit the object partial view. More recently Kroemer et al. [14] proposed an extrusion-based completion approach that is able to deal with shapes that symmetry-based methods cannot handle. The method starts by detecting potential planes of symmetry by combining the Thrun and Wegbreit method with Mitra et al.'s fast voting scheme [15]. Given a symmetry plane, an ICP algorithm is used to decide the extrusion transformation to be applied to the object partial point cloud. Despite the fact that these methods were shown to be robust to noise and were able to deal with a wide range of object classes, they are inherently complex in terms of computational effort and thus, not suitable in real-time. Nevertheless, to simplify this problem, one can take advantage of common scenario structures and objects properties that are usually found in daily environments. They mostly involve man-made objects that are typically symmetric and standing on top of planar surfaces. For example, Bohg et al. [16] took advantage of the table-top assumption and the fact that many objects have a plane of reflection symmetry. Starting from the work of Thrun and Wegbreit [13] and similar in spirit to Bohg et al. [16], we propose a new computationally efficient shape completion approach which translates a set of environmental assumptions into a set of approximations, allowing us to reconstruct the object point cloud in real-time, given a partial view of the object.

B. Task-dependent grasping

Since grasping is highly correlated with the task to be performed on the object, a lot of recent work has focused on incorporating task constraints in robot grasping. This is mostly done by learning a direct mapping function between good grasps and geometrical and action constraints, action features and object attributes. A part of this work focuses on Bayesian network learning to integrate symbolic task goals and low-level continuous features such as object attributes, action properties and constraint features [17], [18]. The goal is to learn features of importance for grasping knowledge transfer. This work is extended to consider object categorical information as an additional feature to predict suitable task-dependent grasping constraints [19]. Further, Detry et al. [20], [21], [22] identify grasp-predicting prototypical parts by which objects are usually grasped in similar ways. The discrete part-based representation allows robust grasping. Differently, in addition to the semantic parts, we also consider a task-dependent setting that uses probabilistic logic and world-knowledge to reason about best pre-grasps. Several approaches make use of object affordances for grasping. While in [23] the authors employ estimated visual-based latent affordances, the work in [24] reasons about grasp selection by modeling affordance relations between objects, actions and effects using either a fully probabilistic setting or a rule-based ontology. In contrast, we employ a probabilistic logic-based approach which can generalize over similar object parts and several object categories and tasks.

Related to our probabilistic logic pipeline is the fully probabilistic one introduced in [25]. It combines low-level features and Bayesian networks to obtain possible task-dependent grasps. Closely related is the semantical pipeline presented in [26]. It employs a semantic affordance map which relates gripper approach directions to particular tasks. However, we exploit additional object/task ontologies using probabilistic reasoning and leverage low-level learning and semantic reasoning. This allows us to experiment with a wide range of object categories.

C. SRL for robot grasping and other robotic tasks

From a different point of view, probabilistic relational robotics is an emerging area within robotics. Building on statistical relational learning (SRL) and probabilistic robotics, it aims at endowing robots with a new level of robustness in real-world situations. We review some recent successful contributions of SRL to various robotic tasks. Probabilistic relational models have been used to integrate common-sense knowledge about the structure of the world to successfully accomplish search tasks in an efficient and reliable goal-directed manner [27]. Further, relational dependency networks have been exploited to learn statistical models of procedural task knowledge, using declarative structure capturing abstract knowledge about the task [28]. The benefits of task abstraction were shown in [29], where the robot uses vague descriptions of objects, locations, and actions in combination with the belief state of a knowledge base for reasoning. The goal of this work is to robustly solve the planning task in a generalized pick and place scenario. Abstract knowledge representation and symbolic knowledge processing for formulating control decisions as inference tasks have proven powerful in autonomous robot control [30]. These decisions are sent as queries to a knowledge base. SRL techniques using Markov Logic Networks and Bayesian Logic Networks for object categorization from 3D data have been proposed in [31].

In probabilistic planning, relational rules have been exploited for efficient and flexible decision-theoretic planning [32] and probabilistic inference has proven successful for integrating motor control, planning, grasping and high-level reasoning [33]. In mobile robotics, relational navigation policies have been learned from example paths with relational Markov decision Processes [34]. In order to compute plans comprising sequences of actions and in turn be able to solve complex manipulation tasks, reasoning about actions on a symbolic level is incorporated into robot learning from demonstrations [35]. Symbolic reasoning enables the robot to solve tasks that are more complex than the individual, demonstrated actions. In [36] meaningful symbolic relational representations are

used to solve sequential manipulation tasks in a goal-directed manner via active relational reinforcement learning. Relational Markov networks have been extended to build relational object maps for mobile robots in order to enable reasoning about hierarchies of objects and spatial relationships amongst them [37].Related work for generalizing over doors and handles using SRL has been proposed in [38].

All of these approaches successfully intertwine relational reasoning and learning in robotics. However, none of these frameworks solves the generalization capability needed for task-dependent grasping following an affordance-based behavior. Relational affordance models for robots have been learned in a multi-object manipulation task context [39]. We propose a probabilistic logic pipeline to infer pre-grasp configurations using object-task affordances. Dealing with unknown objects is an important research topic in the field of robotics. This is because in applications related to object grasping and manipulation, robots aimed at working in daily environments have to interact with many never-seen-before objects and increasingly complex scenarios.

III. Vision-based Scene Description

The role of the visual module (cf. first module box in Fig. 2) is to obtain a semantic description of the perceived objects in terms of their pose, symbolic parts and probability distributions over possible object categories. The object segmentation step [40] is followed by part detection and object category estimation, which rely on a full object point cloud. When only a partial view of the object is available, we employ a symmetry-based methodology for object shape completion. Next, the extraction of semantical parts is based on the object's dimensions along the main geometrical axes and can be achieved by bounding-box analysis via PCA. The low dimensional and efficient representation obtained guides the division of each object into a set of semantical parts, namely, *top*, *middle*, *bottom*, *handle* and *usable area*. This reduces the search space for robot grasp generation, prediction and planning. The next subsections explain our symmetry-based method for shape-completion and the division of the completed point cloud into a set of semantical parts.

A. Object shape completion

As any other type of reconstruction based on single views, computing the bounding-box of the object is an ill posed problem due to lack of observability of the self-occluded part. Thus, as for grasping procedures it is necessary that the robot knows the complete shape of the object of interest, some assumptions about the occluded part must be made. Inspired by the work of Thrun and Wegbreit [13] and Bohg et al. [16] and with computational efficiency in mind, we propose a new approach that translates a set of assumptions and rules of thumb observed in many daily environments into a set of heuristics and approximations. They allow us to reconstruct the unobserved part of an object point cloud in real-time, given a partial view.

We consider the following assumptions: a) the objects stand on top of a planar surface (table top); b) the camera is at a higher viewpoint; c) the objects have rotational symmetry; d) their main geometrical axes are either orthogonal or parallel to the supporting plane; e) the axis of symmetry corresponds to one of the main geometrical axes; and f) the direction of the axis of symmetry indicates the object's pose (i.e., upright or sideways).

Fig. 3: Objects having rotational symmetry.

These constraints model perfectly simple box-like and cylinder-like object shapes, such as kitchen-ware tools, and are reasonable assumptions for many other approximately symmetric objects, such as tools (see Fig. 3). Analogous to [16], we consider only one type of symmetry, however we employ the *line reflection symmetry* [13] as it copes better with the object categories that we want to detect.

Let $P = \{\mathbf{p}\} \subset \mathbb{R}^3$ be the set of visible object surface points. Our shape completion algorithm finds the object symmetry axis s reflecting all visible surface points across it. This corresponds to rotating P around s by 180°. We determine s by analyzing the box that encloses the set P, considering the principal directions of the box and the dimensions along those directions. The symmetry axis (i.e., principal direction) is orthogonal to the cross product of the bounding box directions whose dimensions are the closest, and passes through the bounding-box centroid. To cope with the supporting plane assumption, we compute the horizontal (i.e., table plane, xy) and vertical (i.e. table normal, z) bounding-box directions and their dimensions separately. The vertical direction of the bounding-box is given by the normal vector to the table plane and its length is given by the furthest point from the supporting plane $d_z = \max_z(P)$. Since the horizontal directions are arbitrarily oriented in the supporting plane, we apply the projection of P onto the table plane and compute the directions and their dimensions in that space. The 2D components of the centroid location on the table plane cannot be correctly estimated from a partial view in most of the cases (as illustrated in Fig. 4). Let $W = \{\mathbf{p}_{z=0}\} \subset \mathbb{R}^2$ be the set containing the projected points. We assume that the top part of the object is visible and holds the symmetry assumptions so that the object's xy-centroid, c_{xy}, is obtained by considering only the top region points $W_{\text{top}} = \{\mathbf{w}_{\text{top}}\} \subset W$, satisfying the condition:

$$\mathbf{w}^i_{\text{top}} = \begin{cases} \mathbf{p}^i_{z=0} & \text{if } p_z > \sigma d_z \\ \emptyset & \text{otherwise} \end{cases} \quad (1)$$

where $\sigma \in [0,1]$ is a parameter tuned according to the camera view-point and the object shape curvature (i.e., σ is higher for cylinder-like shapes and lower for parallelepiped-like ones). The eigenvectors provided by PCA on the set W define

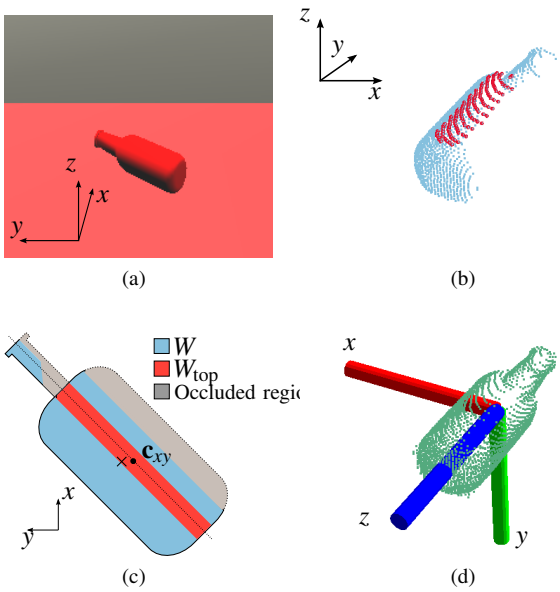

Fig. 4: 2D centroid estimation in the presence of self-occlusion. (a) Bottle camera-view. (b) Visible region (blue) and top visible region surface points (red). (c) Bottle planar projection: $\times$ marks the centroid of W (blue), whereas $\bullet$ indicates the centroid of W_{top} (red). (d) After shape completion, an object coordinate frame is defined as having its origin at the bounding box centroid and z-axis aligned with the symmetry axis.

the horizontal directions whereas their lengths are given by projecting the points in W onto its eigenvectors and finding the maximum in each direction.

B. Part-based object representation

We consider two main types of objects: tools and other objects. A tool has as parts a *handle* and a *usable area*, while the rest of the objects have *top*, *middle*, *bottom* parts and may have *handles*. When the axis of symmetry is parallel to the supporting plane and the lengths of the remaining directions are smaller than a predefined threshold, we consider that the object has a *handle* and a *usable area*. In order to cope with objects such as mugs and pans we detect a handle if a circle is fitted in the projected points W with a large confidence. The points lying outside of the circle are labeled as *handle*. The rest of the points are divided along the axis of symmetry into *top*, *middle* and *bottom*. Fig. 5 illustrates examples of detected semantic parts for several objects using our completion algorithm.

The bounding boxes of the object parts define the pre-grasp hypotheses, providing two pre-grasp poses for each face of a box, as illustrated in Fig. 6. The final number of pre-grasp hypotheses is pruned in a first stage by the task-dependent logical module and in a second stage by a collision checker and the motion trajectory planner.

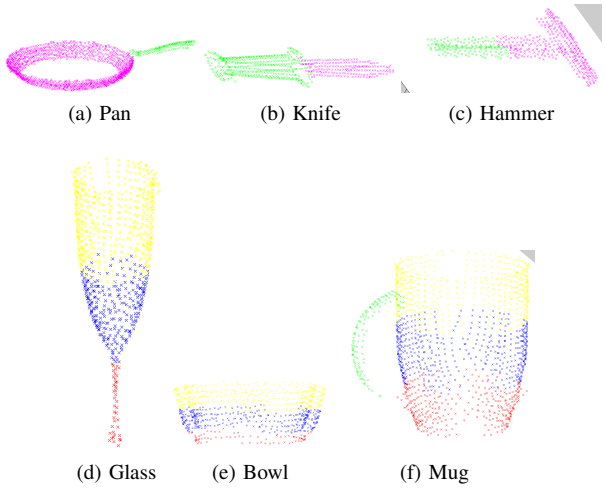

Fig. 5: Semantic parts for several objects after applying the completion algorithm. The colors correspond to parts as follows: yellow - top, blue - middle, red - bottom, green - handle, and magenta - usable area.

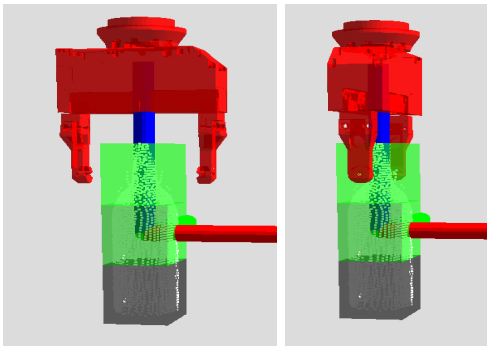

Fig. 6: Examples of the pre-grasp gripper poses for a face of the top part of a bottle.

C. Object category via manifolds

Given the completed object point cloud and its semantic parts, we estimate the object category by retrieving the objects with the most similar global properties. This manifold- and part-based category classification approach was introduced in [41]. While beneficial for direct grasping point prediction [6], global object similarity also ensures a strong enough appearance-based predictor for object category. This prediction in the form of a distribution on object categories is used as a prior for the probabilistic logic module. To incorporate global object similarity for object category prediction, we leverage propagation kernels, a recently introduced graph kernel designed for classification and retrieval of partially labeled graphs [42].

We obtain the distribution on object categories for a particular query object by retrieving the objects in an object database being most similar in terms of global shape and semantic part information. We represent the objects by labeled graphs, where the labels are the semantic part labels derived by the visual module and the graph structure given by a k-nearest neighbor

(k-nn) graph. For each object point cloud we derive a weighted k-nn graph by connecting the k nearest points w.r.t. Euclidean distance in 3D. We use a four-neighborhood (i.e., $k = 4$) and assign an edge weight reflecting the tangent plane orientations of its incident nodes to encode changes in the object surface. The weight of edge (i, j) between two nodes is given by $w_{i,j} = |\mathbf{n}_i \cdot \mathbf{n}_j|$, where $\mathbf{n}_i$ is the normal of point i. The nodes have five semantic classes encoding object part information *top*, *middle*, *bottom*, *handle* and *usable area*. To be able to capture manifold information as graph features in presence of full label information we use a diffusion scheme of the labels corresponding to the *diffusion graph kernel*, in the following simply referred to as *propagation kernel*, proposed in [42]. We stress that the graphs of the 3D point clouds as illustrated[1] in Fig. 1③ capture both manifold information (geodesic distance) via their structure and semantic information (part labels) via their node labels.

The similarity measure among objects is a kernel function over counts of similar node label distributions per diffusion iteration. The T-iteration propagation kernel between two graphs G' and G'' is, then, defined as:

$$K_T(G', G'') = \sum_{t=0}^{T} ker(G'_t, G''_t), \tag{2}$$

where T represents the maximum number of label propagation interactions considered and *ker* is a linear base kernel, defined as:

$$ker(G'_t, G''_t) = \langle \phi(G'_t), \phi(G''_t) \rangle. \tag{3}$$

In our experiments we vary $T \in \{0, ..., 15\}$ and use the maximum number of iterations giving the best results.

The main ingredients of propagation kernels are the distribution-based graph features $\phi(G_t)$. They are essentially computed from node label distributions of running label diffusion on the respective graphs. Hence, the node label distributions of G_t are updated according to $L_t \leftarrow T L_{t-1}$, where the transition matrix T is the row-normalized adjacency matrix $T = D^{-1}A$ and D being the diagonal degree matrix with $D_{ii} = \sum_j A_{ij}$. Based on the $L_t (0 \leq t \leq T)$, we compute for each graph the counts of similar distributions among the respective graphs' nodes. As the node label distributions are $m-$dimensional *continuous* vectors, where m is number of semantic labels, in this case $m = 5$, propagation kernels use locality sensitive hashing (LSH) [43] as a quantization function to ensure the acquisition of meaningful features. We employ a quantization function for distributions preserving the total variation distance (for more details, see [42]). The bin width parameter of LSH is fixed to $w = 10^{-4}$ in all experiments.

Propagation kernels leverage the power of evolving continuous node label distributions as graph features and hence, capture both manifold information and semantic information. Given a new object G^* that the robot aims to grasp, we first select the top n most similar graphs $\{G^{(1)}, \cdots, G^{(n)}\}$, where the similarity is given by the respective row of the correlation matrix $\hat{K}$, where $\hat{K}_{ij} = \frac{K_{ij}}{\sqrt{K_{ii} K_{jj}}}$. Note, that by using $\hat{K}$ instead of K we achieve a normalization w.r.t. the number of points in the point clouds and hence, w.r.t. the scale of the objects. We set $n = 10$ in all our experiments. Second, we build a weighted average over the categories of the objects corresponding to $\{G^{(1)}, \cdots, G^{(n)}\}$ where the weight function is defined as $f(x) = \exp(-x)$ with x being the rank after sorting the kernel row $\hat{K}_{*,:}$. This average is finally used as a prior distribution on the object category for the object with the graph representation G^*. The prior distribution over object categories for the cup in Fig. 1 is:

```
0.56 cup; 0.36 can; 0.05 pot; 0.02 pan
```

and will be used further by the probabilistic logic module to reason about the different prediction tasks.

[1]The illustration does not depict the edge weights being proportional to the change in curvature of the adjacent points.

IV. The Probabilistic Logic Module

After the visual module, we introduce our reasoning module (cf. second module in Fig. 2). Its role is to answers three types of queries, related to object category, most affordable task and best semantic pre-grasp. For example, for object category prediction we query the object instance for being a hammer by calculating the probability $P(cat(O, hammer)|obs, M)$, where O is the object to classify, M is the model and *obs* is the conjunction of observations made about the world (e.g., *obs* = {parts, pose, task}). Similarly, we query for the most likely pre-grasp by calculating probabilities $P(grasp(Pt)|obs, M)$, where Pt is an object part. When the task is not observed the set of observations becomes *obs* = {parts, pose} and then the task can be also inferred from the model by calculating probabilities $P(affords(O, T)|obs, M)$, where T is a task.

Bayesian Networks (BNs) are often used to model such complex dependencies involving uncertainty [19]. Differently, our probabilistic model is defined using Causal Probabilistic (CP) Logic [44]. There are several advantages of a CP-Logic. First, it can intuitively integrate world knowledge as logic rules. For example, we can exploit object ontologies to reason about object (super-)categories. Similarly, we can use object/task ontologies and task-object affordance models to reason about possible, impossible and desirable pre-grasps. Second, CP-Logic is designed to explicitly model causal events (or relationships between random variables). For example, if the object has a usable area and a handle it is likely to be a 'tool' and it can be one (any) of the tool sub-categories (e.g., 'hammer'). This rule is a general, but local piece of information which does not consider other possible causes for the object sub-category. This is rather difficult to encode with a BN, as querying for $P(cat(O, hammer)|obs, M)$ involves knowing all the possible causes for $cat(O, hammer)$ and how they interact with the observations *obs*. Similarly, if the object is a 'tool' and the task is 'pass', then it should be rather grasped by the usable area instead of by the handle. This involves again local causation. In fact, robotic grasping is characterized by a number of causal uncertain events which sometimes involve different consequences. Third, a CP-theory is more efficient as it requires fewer parameters [45] and allows parameter sharing by generalizing over similar situations.

Our grasping CP-theory has 4 parts: 1) a set of probabilistic observations about the world consisting of visual object properties and, optionally, a probability distribution on the task

type; 2) a set of logic rules, defined as background knowledge which incorporate common sense knowledge about the world, that is, object/task ontologies and object-task affordances; 3) a set of probabilistic logic rules as the object category model; and 4) a task-dependent grasping model in the form of probabilistic logic rules. Note that for all prediction tasks we make the mutually exclusiveness assumption. For object category prediction this implies that an object cannot have several categories at the same time. Similarly, for task selection, this translates into the fact that only one task can be executed at any point in time. We use a ProbLog implementation [46] of the CP-Logic theory and we show experimentally that by putting together probabilistic and logical reasoning we improve the grasping performance. We expose the numerated parts in the following subsections.

A. Observations about the world

In this paper, we observe one object at a time.[2] *Visual and task observations* of the scenario in Fig. 1 are shown below. They can be ground probabilistic facts, such as $0.8 :: part(top,o)$ stating that the object o has a top part with probability 0.8, deterministic facts, such as $empty(o)$ stating that o is empty, or CP-rules. A CP-rule capturing the prior distribution over the object category for $object(o)$ is:

```
0.56::cup(o); 0.36::can(o); 0.05::pot(o); 0.02::pan(o)<-
   object(o).
```

stating that an object o belongs to a category with a certain probability, that is, it is either a cup with probability 0.56 or a can with probability 0.36 or a pot with probability 0.05 or a pan with probability 0.02. In any CP-rule, the sum of the possible outcomes can be at most 1.[3] The probability should be interpreted as: if the body is true, then it causes the consequence to become true with a causal probability. Similarly, we can have a prior on the task type as a CP-event. In our experiments, if the task is not given, we assume a uniform distribution on the task type. Similarly, if the prior on the object category is not observed, we consider a uniform prior instead.

```
%visual observations for object o
object(o).
0.8::part(top,o). % the prob. that o has a top is 0.8
1.0::part(handle,o).
1.0::part(middle,o).
1.0::part(bottom,o).
0.5::pose(o,upright).
empty(o).
0.56::cup(o); 0.36::can(o); 0.05::pot(o); 0.02::pan(o)<-
   object(o).

%all tasks t1,...,t7 observed
pourOut(t1).
pass(t2).  ...
%prior on the task type
1/7::task(t1); 1/7::task(t2); ... ; 1/7::task(t7) <- true.
```

CP-rules with probability 1.0 are encoded deterministically.

B. World knowledge: ontologies and affordances

The object ontology is illustrated in Fig. 7 and structures 11 object categories that we consider in our scenario: $C = \{$*pan, pot, cup, glass, bowl, bottle, can, hammer, knife, screwdriver, cooking_tool*$\}$. The super-categories, defined based on the object functionality, are: *kitchenContainer, dish, openContainer, canister, container, tool, object*. By making use of the ontology structure, the grasping model makes abstraction of the fine-grained object categories.

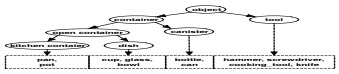

Fig. 7: Object Ontology

Fig. 8: Task Ontology

The task ontology in Fig. 8 structures 7 tasks: $T = \{$*pass, pourOut, pourIn, p&pInUpright, p&pInUpsidedown, p&pInSideways, p&pOn*$\}$. For example, the task *pourOut* refers to the action of removing the contained liquid, while *p&pInUpsidedown* refers to picking and placing the object inside a shelf in the upside-down pose. Depending on the object properties, its parts and the task to be performed, the object should be grasped in different ways.

The object-task affordances for the considered scenario are illustrated in Fig. 9. They allow us to relate the two concepts and thus define the grasping model in a relational way. Both the affordances table and the object ontology were defined by human experience and inspired by *AfNet: The Affordance Network*[4]. They can be extended to include new object/task categories.

This knowledge is translated into deterministic logical rules[5] in the following way:

```
%Examples of object ontology mappings
canister(X):- can(X).
dish(X):- cup(X).
tool(X):- hammer(X).
container(X):- canister(X).
object(X):- tool(X).
object(X):- container(X).
% Examples of task ontology mappings
pour(T):- pourIn(T).
```

[2]The work is easily extendable to consider several objects simultaneously.

[3]If the sum is less than 1, there is a non-zero probability that nothing happens.

[4]Available at: www.theaffordances.net.

[5]For convenience, in the case of task prediction, we use the affordance rules with a high equal probability due to the mutually exclusiveness constraint.

Fig. 9: Object-Task Affordances

```
pour(T):- pourOut(T).
task(T):- pour(T).
% Task affordances
possible(X,T):- object(X), pass(T).
possible(X,T):- container(X), pour(T).
possible(X,T):- object(X), p&pOn(T).
possible(X,T):- container(X), p&pIn(T).
possible(X,T):- tool(X), p&pInSideways(T).
% Some impossible task affordances
impossible(X,T):- canister(X), pourIn(T).
impossible(X,T):- kitchenContainer(X), pourOut(T).
% Common sense exceptions
impossible(X,T):- pan(X), full(X), pass(T).
impossible(X,T):- container(X), full(X), p&pInUpsidedown(T).
...
affords(X,T):- possible(X,T), not(impossible(X,T)).
```

where, for example, $dish(X) : -cup(X)$ is a deterministic intensional rule stating that "any cup is a dish"; $dish(X)$ is the head of the rule, while $cup(X)$ is the body. We initially assume a deterministic affordances model. However, the model can be learned to obtain better estimations [39] or can be re-estimated in our reasoning module by inferring the most likely task. This shows the flexibility of our approach.

C. The CP-theory for semantic grasping

Similarly, we use deterministic rules and CP-events indicating object category consequences based on the object parts and properties. For example the deterministic rule:

```
tool(X):- part(ua,X),part(ha,X),pose(X,sideways).
```

reads as: if the object has a usable area and a handle and it poses sideways, then it is a tool. When observed parts are, for example, *bottom*, *middle* and *top*, no handles are detected and the pose is sideways or upright, then the object can be a glass, a bowl or a canister. If the observed pose is upsidedown then the object can be a glass, a bowl or a can. If exactly one handle is observed, then the object may be a cup or a pan. In these cases we can define CP-events showing the possible outcomes:

```
0.25::glass(X); 0.25::bowl(X); 0.5::canister(X) <-
    part(top,X), part(middle,X), part(bottom,X),
    no_handle(X), pose(X,upright).
0.25::glass(X); 0.25::bowl(X); 0.5::canister(X) <-
    part(top,X), part(middle,X), part(bottom,X),
    no_handle(X), pose(X,sideways).
0.33::glass(X) ; 0.33::bowl(X) ; 0.33::can(X) <-
    part(top,X), part(middle,X), part(bottom,X),
    no_handle(X), pose(X,upsidedown).
0.75::cup(X); 0.25::pan(X) <- part(middle,X),
    part(top,X), part(bottom,X), part(handle,X),
    pose(X,upright).
```

These rules encode generality also by using object super-category atoms in the head. Thus, in order to estimate the object category, we replace the original object ontology defined in Section IV-B with an ontology that models category distributions with respect to the super-categories across the ontology. This is part of the categorization model and is done using CP-events. The causal probabilities are estimated based on the number of specific categories in the leafs. For example, we have the distribution over *hammer*, *knife*, *screwdriver* and *kitchen_tool* caused by the super-category *tool* or the distribution over *can* and *bottle* caused by the object being a *canister*:

```
0.25::hammer(X); 0.25::knife(X); 0.25::screwdriver(X);
0.25::kitchen_tool(X) <- tool(X).
0.5::can(X); 0.5::bottle(X) <- canister(X).
```

Thus, in our experiments, the object categorization CP-theory M contains category rules, world knowledge and visual observations. The CP-theory parameters should not be interpreted as the conditional probability of the head atom given the body, e.g., $P(can(o)|canister(o)) = 0.5$ is incorrect. It is part of the semantics of CP-Logic that each rule independently makes a head atom true when triggered. Thus, the conditional probability that o is a can, given that o is a canister, may be different than 0.5, in case there is a second possible cause, e.g., a prior knowledge that o is a can with probability 0.36, which contributes to $P(can(o))$.

Querying for the most likely object category is equivalent to calculating $argmax_C P(cat(o,C)|M)$. It is possible to ask the query without grounding the specific object category. This will result in a probability distribution over object categories, which is better than the observed prior. For the example in Fig. 1 the new distribution is $P(cat(o,cup)) = 0.98, P(cat(o,pan)) = 0.02$, while $P(cat(o,can))$ and $P(cat(o,pot))$ become 0.

There are different levels of generalization with respect to the rules of the theory. We experimented also with more general rules to investigate the suitability of our model. A more general theory was also able to improve the object category prior (see section VI), showing similar behavior and results to the more specific one. Examples of more general rules are shown below, where we replace, for example, the more specific head *0.25::glass(X); 0.25::bowl(X); 0.5::canister(X)* with the super-category *1.0::container(X)*, while keeping the same rule body. For the example in Fig. 1 the new distribution with the more general theory becomes *P(cat(o,cup))=0.93, P(cat(o,pot))=0.05, P(cat(o,pan))=0.02*.

```
1.0::container(X) <- part(middle,X), part(top,X),
    part(bottom,X), no_handle(X), pose(X,upright).
0.6::dish(X); 0.4::canister(X) <- part(top,X), part(middle,X),
    part(bottom,X),no_handle(X), pose(X,sideways).
1.0::container(X) <- part(middle,X), part(top,X),
    part(bottom,X), no_handle(X), pose(X,upsidedown).
0.5::cup(X); 0.5::kitchen_container(X) <- part(middle,X),
    part(top,X), part(bottom,X), one_handle(X).
```

To *query for the most likely task* in a grasping scenario, we use world observations, the probability distribution over object categories and world knowledge. We can estimate the most likely task by calculating $argmax_T P(affords(o,T)|M)$, where M is the CP-theory. Again, one can ask queries without grounding T to obtain

a probability distribution over possible task types. For the example in Fig. 1 the distribution over possible tasks is *P(affords(o,pass))=0.32, P(affords(o,p&pOn))=0.32*, *P(affords(o,p&pInUpright))=0.32*,*P(affords(o,pourIn))=0.03*, *P(affords(o,p&pInUpsidedown))=0.01*.

We define the part-based grasping model as a set of CP-events. Each causal event generates as consequence the graspability of a certain object part conditioned on the part existence, task, object (super-)category and properties. The feasibility of the semantic grasp is encoded via the causal probability. Some examples from the grasping model for the *dish* super-category are shown below:

```
0.8::grasp(X,T,middle) <- affords(X,T), pass(T), dish(X),
    pose(X,upright), full(X), part(middle,X).
0.1::grasp(X,T,top) <- affords(X,T), pass(T), dish(X),
    pose(X,upright), full(X), part(top,X).

0.1::grasp(X,T,bottom) <- affords(X,T), pass(T), dish(X),
    pose(X,upright), empty(X), part(bottom,X).
0.6::grasp(X,T,middle) <- affords(X,T), pass(T), dish(X),
    pose(X,upright), empty(X), part(middle,X).
0.2::grasp(X,T,top) <- affords(X,T), pass(T), dish(X),
    pose(X,upright), empty(X), part(top,X).
0.1::grasp(X,T,handle) <- affords(X,T), pass(T), dish(X),
    pose(X,upright), full(X), part(handle,X).

0.2::grasp(X,T,bottom) <- affords(X,T), pass(T), dish(X),
    pose(X,upsidedown), part(bottom,X).
0.7::grasp(X,T,middle) <- affords(X,T), pass(T), dish(X),
    pose(X,upsidedown), part(middle,X).
0.1::grasp(X,T,top) <- affords(X,T), pass(T), dish(X),
    pose(X,upsidedown), part(top,X).
...

0.7::grasp(X,T,middle) <- affords(X,T), p&pIn(T), dish(X),
   pose(X,sideways), part(Id,middle,X).
0.3::grasp(X,T,bottom) <- affords(X,T), p&pIn(T), dish(X),
   pose(X,sideways), part(bottom,X).

1.0::grasp(X,T,middle) <- affords(X,T), pourIn(T), dish(X),
   empty(X), part(middle,X).
1.0::grasp(X,T,middle) <- affords(X,T), pourOut(T), dish(X),
   not(empty(X)), part(middle,X).
...
```

We can enforce constraints to model impossible pre-grasps. For example, when the object is a *tool* and the task is *pour*, we have an impossible affordance and thus, an impossible pre-grasp position. Examples of such constraint rules in ProbLog are:

```
%constraint for impossible affordances
false:- grasp(X,T,R), task(T), object(X),
   impossible(X,T),part(R,X).
%constraint for collision
false:- grasp(X,T,R), task(T), object(X),
   part(R,X), collision(R).
%other constraints
false:- grasp(X,T,R), pose(X,upsidedown),
   pan(X), task(T), part(R,X).
...
```

The first constraint states that it is impossible that the pre-grasp atom $grasp(X,T,R)$ is true when the body is true. This will guarantee that the probability of such grasps is equal to 0. The second constraint rule shows that, additionally, we can connect the reasoning module to the execution planner by enforcing the probability of a pre-grasp to 0 if there are environmental constraints for the gripper. The third constraint indicates that if the object is a pan in the upside down pose then no task should be executed, as grasping the object in this situation is very difficult.

If M is the CP-theory for task-dependent grasping, we can *query for the most likely semantic pre-grasp* of an object. This is equivalent to calculating $argmax_{Pt} P(grasp(o,t_2,Pt)|M)$, where Pt is a part in the set of observed object parts and t_2 is the given task. For the example in Fig. 1 the distribution over possible parts when the task considered is *pass* becomes: *P(grasp(o,pass,middle))=0.87*, *P(grasp(o,pass,top))=0.08*, *P(grasp(o,pass,bottom))=0.03*, *P(grasp(o,pass,handle))=0.01*.

Similar to the object categorization theory, there are different levels of generalization with respect to the rules. To test the fittingness of the theory we experimented also with more general rules, by generalizing over the object pose and containment with respect to several tasks and thus, reducing the number of rules. For example, we replaced part of the theory presented above for task *pass* and super-category *dish* with:

```
0.1::grasp(X,T,bottom) <-
   affords(X,T), pass(T), dish(X), part(bottom,X).
0.6::grasp(X,T,middle) <-
   affords(X,T), pass(T), dish(X), part(middle,X).
0.2::grasp(X,T,top) <-
   affords(X,T), pass(T), dish(X), part(top,X).
0.1::grasp(X,T,handle) <-
   affords(X,T), pass(T), dish(X), part(handle,X).
```

We have defined our models using human experience and "educated guesses". They can be augmented by adding extra rules to include new object/task categories. The world knowledge was encoded as general as possible while still reflecting the ontologies and task-object affordances. The parameters of the rules composing the models can, in principle, be learned from data [45] to best represent the application domain. Our current experimental results with the quite rigid affordance model can be improved by learning better probability estimates for object-task affordances from data.

V. Local Shape Grasping Prediction

The probabilistic logic module selects the object part to be grasped and/or the task to be performed by taking into account high level information. Starting from here, the third module of the pipeline (cf. third module box in Fig. 2) uses additional local shape features characterizing the object part to compute the grasping probability. It calculates $P(grasp|local\ shape)$ by mapping the classification output of a Support Vector Machine (SVM) onto a probability. The SVM classifier discriminates between graspable and non-graspable shapes.

A. Depth difference features

The local shape features are computed in the volume enclosed by the gripper, which is a bounding box located and oriented according to the pre-grasp hypothesis pose. Depth changes in the objects were shown helpful to recognize graspable regions, even in cluttered environments where objects cannot be segmented accurately [47]. The symmetry height accumulated feature [47] is robust, but constrained to top grasps only. We introduce a feature with computations based also on heights, however, it can be computed for any grasping orientation. Our feature, called *depth gradient image* (DGI), computes the gradient of the depth image in the volume enclosed by the gripper. This volume defines a depth value (i.e., the height in mm) as the z-component of the distance

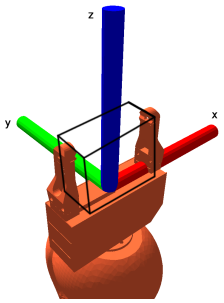

Fig. 10: Gripper and volume of interest, showing the reference frame origin for the orthogonal projection of the DI image from Eq. (4).

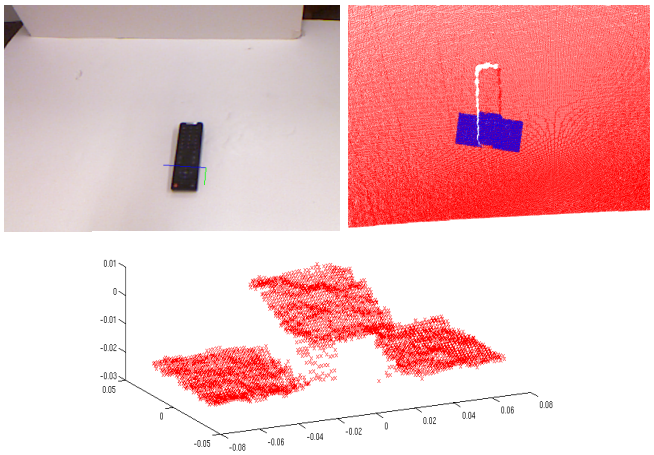

Fig. 11: Object and its correspondent point cloud and cropped grasping cloud. The top right and bottom images show the volume enclosed by the gripper of an object. The blue points show the selected points of a graspable region of the remote control. The bottom image shows the points enclosed by the gripper volume.

from the gripper base to the object point. Fig. 11 shows an example of the selected region of an object and Fig. 10 illustrates the volume of interest enclosed by the gripper.

The depth image requires a discrete sampling of the volume, which was defined as boxes of 7x7x15 (mm) and is defined as:

$$\mathrm{DI}(u,v) = \begin{cases} \min\{z\} & \text{if } z \in \mathrm{box}(u,v) \\ -1 & \text{otherwise,} \end{cases} \tag{4}$$

where $\mathrm{box}(u,v)$ represents the set of points inside the box defined at the pixel (u,v). Eq. (4) performs an orthogonal projection of the closest point to the base of the gripper for every pixel of the depth image. Fig. 12 shows the depth image for the selected volume in Fig. 10. Finally, the DGI is computed on the depth image by applying pixel differences in u and v as follows:

$$\mathrm{DI}_u(u,v) = \mathrm{DI}(u+1,v) - \mathrm{DI}(u-1,v), \tag{5}$$

$$\mathrm{DI}_v(u,v) = \mathrm{DI}(u,v+1) - \mathrm{DI}(u,v-1), \tag{6}$$

$$\mathrm{DGI}(u,v) = \sqrt{\mathrm{DI}_u(u,v)^2 + \mathrm{DI}_v(u,v)^2}. \tag{7}$$

The DGI acts as a local shape descriptor for grasping prediction. The descriptors of several graspable and non-graspable regions are next fed into a SVM classifier.

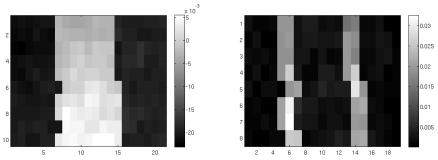

(a) Depth Image (b) Depth Gradient Image

Fig. 12: Example of a depth image (10x21 pixels) and its corresponding gradient magnitude (8x19 pixels).

B. Grasping probability

Given DGI shape features $\mathbf{x}_i$ and their labels y_i, we use SVMs [48] with Radial Basis Function (RBF) kernel to discriminate between successful and failed grasps. Before applying the sign function, we map the SVM output:

$$f(\mathbf{x}) = \mathbf{w}^T \phi(\mathbf{x}) + b \tag{8}$$

$$\phi(\mathbf{x}_i) = \exp(-\gamma|\mathbf{x}_i - \mathbf{x}_j|^2), \tag{9}$$

onto a probability by applying a sigmoid function to the decision value from Eq. (8). We employ the parametric sigmoid:

$$P(grasp|local\ shape) = \frac{1}{1+\exp(Af(\mathbf{x})+b)}, \tag{10}$$

where the parameters A and b are obtained by generating a hold-out set and cross-validation. Its advantages were shown empirically in [49].

C. Learning dataset and evaluation

We train our classifier on the grasping rectangle dataset [50]. It contains camera images, point clouds and the pre-grasp poses for both successful and failed grasps. The dataset contains more than 5K samples with similar proportions between positive and negative samples. We remove objects with very noisy point clouds, so in total we have 4708 samples (2424 positive and 2284 negative). We apply cross-validation with 4 splits in order to find the best RBF parameter ($\gamma = 50$) and sigmoid parameters ($A = 2$ and $b = 0.5$). The grasping accuracy defined in [50] selects the top grasping region per object and then compares it to the ground truth. Table I shows that our approach using only depth features has good performance, improving the result in [50]. A better result is obtained in [3], however they rely on both image and depth features.

Approach	Grasping accuracy
Jiang et. al. [50]	84.7
Lenz et. al. [3]	93.7
this work	92.63

TABLE I: Grasping accuracy (%) on the rectangle dataset [50].

VI. Experiments

We address experimentally the benefits of our proposed pipeline for robot grasping. Specifically, we investigate the following main questions:

(Q1) How good is the visual module?
(Q2) How robust is the probabilistic logic module? How well does it cope with missing information?
(Q3) Does the integration of high-level reasoning (about task and object category) and low-level learning improve the grasping performance upon local shape features?

Since the symbolic visual detection plays an important role for the rest of the pipeline, we first answer question (Q1) by investigating the performance of the object pose and part detection (subsection VI-C1). Next, we evaluate the performance of the object category prediction using the manifold-based graph kernel (subsection VI-C2).

We answer question (Q2) by tackling in subsections VI-C2, VI-C3 and VI-C4 the following questions:

(i) does the probabilistic logic module improve upon manifold information for object category prediction?
(ii) can it predict correctly suitable tasks?
(iii) can it predict the correct pre-grasp region?

We investigate its robustness with respect to object category, task, pre-grasp and grasping point prediction in subsections VI-C2, VI-C3, VI-C4 and VI-C5. We perturb either the visual observations about the world by dropping the prior on the object category, or the CP-theory by keeping the more general rules.

Finally, to answer question (Q3) we first learn a classifier that maps points sampled from full objects to successful grasps using local shape features. Given a new object, we then directly predict the most likely grasp using solely local shape information. This is our local shape-based baseline. We compare the baseline with the pipeline classifier, which maps points from semantic pre-grasp object regions to successful grasps, using similar local shape features. Given a new object, we first predict, using high-level reasoning, the most likely object pre-grasp region. We then use the classifier to predict good grasping points among the set of points in the inferred pre-grasp part (the most appropriate for a given task).

A. Datasets and evaluation scenarios

We consider three types of datasets to quantitatively investigate the robustness and power of generalization of our SRL approach. For the first type, the object point clouds are obtained from 3D meshes and the object parts are manually labeled. In this case the dataset is synthetic and actual grasps are not executed. For the second type, data samples are obtained from the ORCA simulator [51], while for the third type, data samples are obtained from a real robot platform.

1) Synthetic scenario: It considers flawless visual detection of objects from 3D meshes. The object points are distributed uniformly on the object surface according to their size by applying the midpoint surface subdivision technique. Point normals are correctly oriented, the object pose and its parts are manually labeled as well as the object containment. This "perfect scenario" serves as an upper-bound comparison scenario to the more realistic scenarios, allowing an extensive evaluation of the generalization capabilities of the probabilistic logic module. The dataset contains 41 objects belonging to all categories in our ontology and 102 grasping scenarios. We denote this dataset S_{SYN}.

2) ORCA scenarios: The second type of datasets is used to evaluate all the modules and the full pipeline in simulation. We use ORCA, which provides the sensors (laser range camera Asus Xtion PRO [52] and the Universal Gripper WSG 50 [53] force sensor), robotic arm (KUKA LightWeight Robot (LWR) [54]), objects and interface to physics engine (Newton Game Dynamics library [55]) for robot grasping simulation. The other modules are external to ORCA and interfaced both with the simulated and real robot. These modules include: object completion, part and pose detection, global shape similarity, probabilistic logical reasoning modules, local shape grasping prediction and the tree-based motion planner [56] available in the Open Motion Planning Library (OMPL) [57]. Each object is placed on top of a table.

Further, we consider four possible settings. In the first setting, object pose is not estimated but given by the ground truth, while the parts are estimated from the completed point cloud, as explained in Section III. Thus, the scene description may have missing parts when they are occluded or not detected. Additionally, we assign to all detected parts probability 1.0. We denote this dataset S_{REAL_semi}. In the second setting, both object pose and its parts are estimated from the completed point cloud. While the pose has associated a likelihood, we keep highly confident parts. We denote this setting S_{REAL}. In the third setting we also provide a part likelihood according to the limitations of the detection algorithm. We denote this dataset S_{REAL_noisy}. Finally, the fourth setting includes actual grasping tests with the simulated robot. It comprises a subset of the scenarios from the third setting, where all containers are empty and all objects are graspable by the robot. The rationale behind is that it is very difficult to check whether a container is full or if some objects (too big or too small) do not fit the gripper capabilities. We denote this dataset S_{GRASP_noisy}. In addition, object poses considered are *upright* or *sideways* due to the ambiguity between upright and upside-down when using global shape representations. Each of the first three settings contains 26 different objects, instances of categories *pan, bowl, cup, glass, bottle, can, hammer, screwdriver, knife* and 126 grasping scenarios. The fourth setting contains 18 objects, instances of categories *pan, cup, glass, bottle, can, hammer, screwdriver, knife* and 113 grasping scenarios.

3) Real robot scenario: The robot scenario considers the same type of tests as those included in the S_{GRASP_noisy}. Fig. 14 shows our setup. In addition, we evaluate the performance of the pipeline when two or more objects are in the field of view of the camera and the field of action of the arm, considering three settings of increasing complexity in terms of path planning. The less complex setting (scenario1), considers only two objects which are instances of *glass* and *bottle*, in a way that planning constraints are very similar to a single object on the table. The setting with intermediate complexity (scenario2) includes three objects which are instances of *can, hammer* and

Fig. 13: Experimental settings with the real robot. Each picture shows the objects utilized for each scenario. Additional object constraints are: the gray bottle of scenario3 is full with water, the white bottle is empty and the coffee container is full of coffee.

Fig. 14: The table is in front of the mobile platform. The range sensor is marked by the green rectangle.

screwdriver. The more complex setting (scenario3) considers four objects which are instances of *bottle, glass* and *cup*. Fig. 13 shows the objects of every scenario. In addition to the larger number of objects, we also consider object placement as another criterion for evaluation. Object placement is performed in two steps: (i) plan from the grasp pose to a post-grasp pose and (ii) plan from the post-grasp pose to the grasp pose. We denote the associated dataset S_{ROBOT}.

The synthetic dataset and part of ORCA datasets are available for download at http://www.first-mm.eu/data.html.

B. Evaluation measures

We evaluate our experiments in terms of accuracy given by $\frac{\#successes}{\#tries} \cdot 100\%$. We assess a *success* in several ways. Depending on the prediction task, the ground truth is either one value (object categorization, pose detection) or a set of values (task and pre-grasp prediction, part detection).

In the case of object part detection we consider a success if all detected object parts match the ground truth. A correct pose detection is considered when the discrete pose predicted matches the ground truth.

Dataset	Part detection	Pose detection
S_{GRASP_noisy}	84.56	100
S_{ROBOT}	82.14	100

TABLE II: Accuracy (%) for object part and pose detection.

For object categorization we take as prediction the category with the highest probability and consider it a success if it matches the ground truth category. For the uniform prior, it can be that two or more categories are predicted with the same probability. This case is reported as a false positive.

For task/pre-grasp prediction evaluation, the ground truth Gt of each instance is a set (e.g., the tasks *p&pOn* and *p&pInUpright* may be equally possible in a particular scenario). In this case, we compare the set of best predictions Pr to Gt, where $|Pr| \leq |Gt|$. If $Pr \subseteq Gt$ a success is reported. We present results for different sizes of Pr, such that $|Pr|$ belongs to the set $\{|Gt| - i\}$, with i ranging from 0 (the most restrictive evaluation setting) to $|Gt| - 1$ (the most pertinent setting). We denote the possible evaluation settings as E_i. For the scenarios with robot grasping execution (S_{GRASP_noisy} and S_{ROBOT}) the evaluation must consider the success of grasping execution with respect to the valid grasping hypotheses provided by the pipeline. In this case the accuracy is given by $\frac{\#correctly\ estimated\ task/part}{\#valid\ grasping\ hypotheses} \cdot 100\%$.

C. Results and discussion

In the following we present quantitative experimental results for all questions. For all results bold font, if present, indicates the best performance.

1) Object pose and part detection: Results are shown in Table II. We note that the PCA global representation is able to cope well with object pose detection, considering the table-top assumption and the object categories assumed. We note that we do not consider the upside-down pose for these tests, as in real-world applications usual poses are upright and sideways. Object part detection suffers from part occlusion for particular object poses, reducing the pipeline performance for object category prediction. In addition to the accuracy, we stress the execution time for the object completion using symmetries and pose detection. The average execution times are 27.5 ms and 15.71 ms on a PC using one core of the Intel Xeon (2.67GHz). These numbers confirm the computational efficiency of our approach, which allows to make fast decisions.

2) Object category prediction: We next evaluate the object category prediction task and report accuracy results in Table III. We compare random category assignment (Random), a linear kernel among label counts (Label fractions), propagation kernels (Manifolds) and the probabilistic logic module (PLM). Note, that the second baseline method corresponds to setting $T = 0$ in Eq. (2). Hence, we compute a linear base kernel using only label counts as features and no manifold information (graph structure) is considered.

For Manifolds we set the parameter T to give the best leave-one-out accuracy performance $T = 12$ on S_{SYN} and $T = 0$ on

Dataset	Random	Label fractions	Manifolds	PLM$_{uniform}$	PLM$^{general}_{uniform}$	PLM	PLM^{general}
S_{SYN}	9.1	63.4	*87.8*	31.37	31.37	**93.14**	92.16
S_{REAL_semi}	9.1	*39.7*	*39.7*	14.29	14.29	**49.21**	*46.83*
S_{REAL}	9.1	*39.7*	*39.7*	14.29	14.29	**48.41**	*46.83*
S_{REAL_noisy}	9.1	**39.7**	**39.7**	14.29	14.29	**39.7**	**39.7**
S_{GRASP_noisy}	9.1	-	**37.95**	-	-	-	-
S_{ROBOT}	9.1	-	**50**	-	-	-	-

TABLE III: Accuracy (%): PLM vs. propagation kernel (Manifolds) vs. baselines for object categorization.

all S_{REAL}.[6] On S_{SYN} we clearly see that incorporating manifold information, that is, using the graph kernel, improves the accuracy more than 20% upon using label information only. On S_{REAL} the manifold improvement is 4 times more than random.

We also evaluated the Manifolds and baseline approaches in terms of average rank (r) of the true category. We obtain $r = 1.39$ (S_{SYN}) and $r = 2.45$ (S_{REAL}) for Manifolds, $r = 1.88$ (S_{SYN}) and $r = 2.54$ (S_{REAL}) for Label fractions and $r = 5.5$ (S_{SYN}, S_{REAL}) for Random. Note that for the rank lower values are better and the optimal result would be 1, i.e., the true category has the highest probability for all objects in the respective dataset/scenario. On both S_{SYN} and S_{REAL} propagation kernels decrease the average rank compared to Label fractions and Random. The parameter T for Manifolds was set in this case to $T = 13$ on S_{SYN} and $T = 3$ on S_{REAL}. The main problem for predicting the object category in the S_{REAL} scenarios is that the part assignment is achieved by applying the part detector introduced in Section III-B. As this detector was designed to work for general objects (i.e. we assume the scenarios to be as realistic as possible) its performance is rather poor for some object categories. Overall, manifold information leads to a good prior distribution among object categories (more details can be found in [41]).

Next, we incorporated the manifold priors into the PLM. We started from the priors that gave the best accuracy for each task. For object categorization this translates into $T = 13$ for S_{SYN} and $T = 0$ for S_{REAL}. To show the robustness of the PLM, we varied the generality of our categorization theory and experimented with and without the manifold prior on the object category. $PLM^{general}_{uniform}$ indicates the more general theory setting with the uniform prior, while $PLM^{general}$ indicates the more general theory with the manifold prior. The results show that the PLM improves object categorization accuracy upon manifold information. By increasing the generality of the theory, we do not loose much in terms of performance when the manifold information is used, and we are still able to improve upon the prior.

Note that by removing the manifold prior, the PLM still gives a reasonable result (3 times better than Random on S_{SYN}). We also evaluated the PLM$_{uniform}$ with a second accuracy definition (acc) in which a *success* is reported if at least one category in the set of equally and maximum category prediction values is equal to the ground truth. In this case we obtain $acc = 68.25\%$ (S_{REAL} datasets) and $acc = 99.02\%$ (S_{SYN}) for PLM$_{uniform}$, and $acc = 62.70\%$ (S_{REAL} datasets) and $acc = 94.12\%$ (S_{SYN}) for $PLM^{general}_{uniform}$. The results obtained using this evaluation setting explain the good performance for the other grasping tasks in subsections VI-C3 and VI-C4. That is, estimating the category of an object as any of the sub-categories of a super-category in the ontology is satisfactory to predict good semantic pre-grasps.

3) Task prediction: We answer the question *what is the suitable task* by reporting results in Table IV (top rows) with evaluation setting E_0 for both general and more specific object categorization theories. Fig. 15 presents results for the specific object categorization theory using all evaluation settings. We experimented on all datasets with and without the manifold prior. The presence of the prior gives significantly better results in the most restrictive evaluation setting. For the other evaluation settings, in most situations, the probabilistic logic module will return, although not as the first option, a correct possible task with or without a prior.

In the scenarios with grasp execution (bottom rows in Table IV), the evaluation settings E_0 and E_1 consider the outcome of the grasping action. The additional source of failures on grasping include uncertainty on the pose of the objects and the gripper, which are caused by the sensor and the object completion. These sources have effects on the performance of the planner, for instance placing the gripper a bit misaligned or hitting and object before closing the gripper. The few cases where the uniform prior provided better results (S_{ROBOT}) are explained by the fact that the complete pipeline failed on the PLM experiment more than the uniform prior due to uncertainty.

We note the importance of having a prior probability distribution over the object categories, rather than the top category. We perform the same experiments only with the top predicted category and obtain accuracies of 95.24%, 89.68% and 89.68% for S_{REAL_semi}, S_{REAL} and S_{REAL_noisy}, respectively, which are lower than using the full prior (see Table IV).

4) Pre-grasp selection: To answer the question *what is the suitable pre-grasp* we considered the specific object categorization theory and experimented with both a more specific and a more general task-dependent grasping theory. The results for both settings when the task is given for evaluation setting E_0 are shown in Table V. We note that increasing the generality of the model does not cause much performance loss. The result confirms generalization over similar object parts and object/task categories, which implies that if the input object is an unseen category, such as a paint roller or a vase the grasping pipeline is robust enough to return a good grasping

[6]The Manifolds results for S_{REAL_semi}, S_{REAL} and S_{REAL_noisy} are the same due to the fact that object pose and part confidence are not used by the propagation kernel. We refer to these datasets as S_{REAL} for the Manifolds evaluation.

Dataset	PLM$_{uniform}$	PLM$^{general}_{uniform}$	PLM	PLM^{general}
S_{SYN}	71.57	65.69	**72.55**	**72.55**
S_{REAL_semi}	**98.41**	**98.41**	95.24	93.65
S_{REAL}	80.16	80.16	**95.24**	93.65
S_{REAL_noisy}	38.10	38.10	**93.65**	**93.65**
S_{GRASP_noisy} E_0	30.00	-	**35.71**	-
S_{GRASP_noisy} E_1	40.00	-	**50.00**	-
S_{ROBOT} E_0	**28.57**	-	25.00	-
S_{ROBOT} E_1	**85.71**	-	75.00	-

TABLE IV: Accuracy (%): PLM for task prediction.

Dataset	PLM$_{uniform}$	PLM$^{general}_{uniform}$	PLM	PLM^{general}
S_{SYN}	81.23	80.67	**85.29**	84.73
S_{REAL_semi}	69.95	72.00	85.26	**86.73**
S_{REAL}	72.00	72.00	**85.26**	84.69
S_{REAL_noisy}	69.16	69.16	85.49	**86.73**
S_{GRASP_noisy} E_0	66.7	-	**75.51**	-
S_{GRASP_noisy} E_1	66.7	-	**75.51**	-
S_{ROBOT} E_0	**66.7**	-	**66.7**	-
S_{ROBOT} E_1	**66.7**	-	**66.7**	-

TABLE V: Accuracy (%): PLM for pre-grasp prediction.

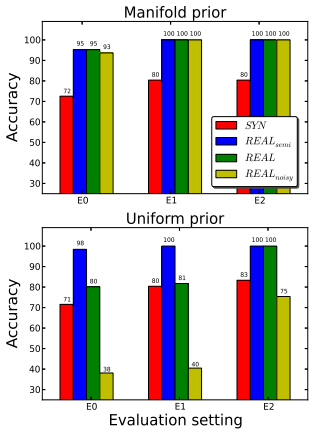

Fig. 15: Accuracy (%): PLM for task prediction using all evaluation settings.

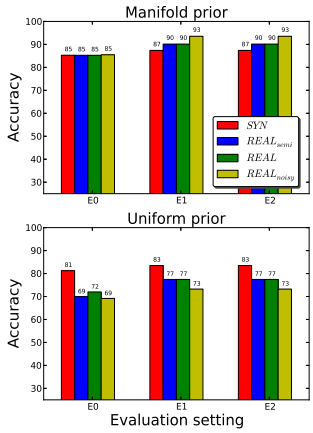

Fig. 16: Accuracy (%): PLM for pre-grasp prediction for all evaluation settings.

part. This allows us to experiment with a wide range of object/task categories and lets us to believe that our approach can be extended beyond the categories used, by augmenting the probabilistic logic module with extra rules.

When the task is predicted by the probabilistic logic module we report accuracies of 95.10% and 100.0% for S_{SYN} and S_{REAL_noisy}, respectively, for both priors using the evaluation setting E_0. When the task is given, results for the pre-grasp prediction task using the most specific task-dependent grasping theory for all evaluation settings are reported in Fig. 16.

For this task as well, in the scenarios with grasp execution, the evaluation considers the outcome of the grasping action. In all cases the object category provided by the manifold matching has better performance than the uniform category prior, showing the benefits of the pipeline. Again, it is important to consider the full prior distribution as input to the PLM, instead of only the top category. Our experiments using solely the top category resulted in accuracies of 81.63, 81.63 and 82.88, respectively, for the three S_{REAL} datasets, which are lower than the ones using the full prior. For both task and pre-grasp prediction we incorporated the manifold priors using parameters $T = 13$ for S_{SYN} and $T = 3$ for S_{REAL}.

The robot experiments (S_{GRASP_noisy} and S_{ROBOT}) did not consider the mutually exclusiveness assumption. However, for the other datasets we performed experiments both with and without this assumption. The means of the differences between the mutually exclusive and the non-exclusive results for object categorization, task prediction and pre-grasp prediction (1.6%, 8.5% and 2.6%, respectively) indicate a better result for the mutually exclusive setup. This lets us to believe that current robot results can be only improved by the exclusiveness assumption.

5) Full pipeline evaluation: The selection of grasping points using only local shape descriptors bias the ranking of the object points towards the most "visually graspable", disregarding other constraints such as the pre-grasp pose for task execution, path planning and post-grasp object pose. By adding those constraints, regions with lower "visually graspable" probability will become more important when considering the task execution probability and vice versa. Thus, in order to answer question (Q3) we compute the percentage of grasping points that have a low "visually graspable" probability from Eq. (10), but still lead to successful grasps when taking into account task constraints. Results are shown in Table VI, having as baseline the local shape-based approach. We also investigate its robustness with respect to grasping point prediction by dropping the prior on the object category (Pipeline$_{uniform}$).

The full pipeline selects in average more points having

Feature (DGI)	Measure	Local	Pipeline	Pipeline$_{uniform}$
S_{GRASP_noisy}	lt 0.5	44.55%	**52.38%**	50%
	lt 0.4	27.27%	**28.57%**	21.15%
	lt 0.3	4.55%	**7.94%**	7.69%
S_{ROBOT}	lt 0.5	37.5%	**50%**	46.15%
	lt 0.4	25%	**42.86%**	23.08%
	lt 0.3	12.5%	**28.57%**	23.08%

TABLE VI: Percentage (%) of successfully graspable points that have "visually graspable" probability less than (lt) 0.3, 0.4 or 0.5: Pipeline vs. local shape grasp prediction.

Scenario	Total tests	Reachable pre-grasps	Grasped objects	Placed objects
scenario1	10	9 (90%)	7/9 (77.8%)	7/9 (77.8%)
scenario2	15	10 (66.7%)	7/10 (70%)	7/10 (70%)
scenario3	20	16 (80%)	8/16 (50%)	8/16 (50%)

TABLE VII: Percentage of successful grasps in the real robot scenarios. Different levels of S_{ROBOT} complexity.

low "visually graspable" probabilities than the other options. This behavior confirms the importance of task constraints on the computation of the grasping probability. We note that the complete pipeline clearly improves upon using the local shape-based approach. This answers affirmatively (Q3). Additionally, the results obtained by Pipeline$_{uniform}$ show again the robustness of the pipeline.

Finally, we present results for the grasp and place action for each level of complexity on the S_{ROBOT} dataset in Table VII. It is important to notice the performance drop of 20% in average when the complexity increases from medium (scenario2) to complex (scenario3). Considering that the multiple object scenario of these experiments is rather simple, one way to improve the performance would be to plan for object displacement (sequence of actions before the pre-grasp pose) before the grasp execution in order to increase the grasping performance on scenario3.

We note that in all our experiments with the pipeline *the result changes if the data is perturbed* for the task, pre-grasp and the grasping point prediction. This highlights the importance of the object categorical information for robot grasping. However, when this information is absent the probabilistic logic theory is robust enough to give reasonable results thanks to the use of world general knowledge. The experiments show the *robustness of the probabilistic logic module when more general rules* are used.

VII. Conclusions

We proposed a new probabilistic logic pipeline which combines high-level reasoning and low-level learning for robot grasping. The high-level reasoning leverages symbolic world knowledge, in the form of object/task ontologies and object-task affordances, object categorical and task-based information. The low-level part is based on learning with shape visual features. The non-trivial realization of high-level knowledge relies on logic, which exploits world knowledge and relations to encode compact grasping models that generalize over similar object parts and object/task categories in a natural way. When combined with probabilistic reasoning, our proposed pipeline shows robustness to the uncertainty in the world and missing information. In addition, our experiments confirm the importance of high-level reasoning and world-knowledge as opposed to using solely local shape information for robot grasping.

As future work we point out three important directions. One is learning the parameters and structure of our logic theory from data. Further, we would like to exploit recent advances in the research on persistent homology [58]. Persistent homology is a tool from topological data analysis performing multi-scale analysis on point clouds identifying clusters, holes, and voids therein. Incorporating such a technique into our pipeline can improve the current part detector. Another direction is planning a sequence of actions in order to fulfill the task-dependent pre-grasp poses. Since planning in presence of multiple objects raises complexity and generalization issues, considering relational planners similar to those in [59] may provide successful plans for pre-grasp tasks.

References

[1] J. Bohg and D. Kragic, "Learning grasping points with shape context," *IEEE Robotics and Automation Society*, vol. 58, no. 4, pp. 362–377, 2010.

[2] L. Montesano and M. Lopes, "Active learning of visual descriptors for grasping using non-parametric smoothed beta distributions," *IEEE Robotics and Automation Society*, vol. 60, no. 3, pp. 452–462, 2012.

[3] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *CoRR*, vol. abs/1301.3592, 2013.

[4] A. Saxena, L. L. S. Wong, and A. Y. Ng, "Learning grasp strategies with partial shape information," in *AAAI*, 2008.

[5] J. Aleotti and S. Caselli, "Part-based robot grasp planning from human demonstration," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 4554–4560.

[6] M. Neumann, R. Garnett, P. Moreno, N. Patricia, and K. Kersting, "Propagation Kernels for Partially Labeled Graphs," in *ICML Workshop on Mining and Learning with Graphs*, Edinburgh, UK, 2012.

[7] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 2, pp. 239–256, 1992.

[8] S. Rusinkiewicz and M. Levoy, "Efficient variants of the icp algorithm," in *International Conference on 3-D Digital Imaging and Modeling*, 2001.

[9] Y. Chen and G. Medioni, "Object modelling by registration of multiple range images," *Image Vision Comput.*, vol. 10, no. 3, pp. 145–155, 1992.

[10] Z. C. Marton, L. Goron, R. B. Rusu, and M. Beetz, "Reconstruction and Verification of 3D Object Models for Grasping," in *Proceedings of the 14th International Symposium on Robotics Research*, 2009.

[11] J. Kuehnle, Z. Xue, M. Stotz, J. Zoellner, A. Verl, and R. Dillmann, "Grasping in depth maps of time-of-flight cameras," in *International Workshop on Robotic and Sensors Environments*, 2008, pp. 132–137.

[12] G. Biegelbauer and M. Vincze, "Efficient 3d object detection by fitting superquadrics to range image data for robot's object manipulation," in *IEEE International Conference on Robotics and Automation*, 2007, pp. 1086–1091.

[13] S. Thrun and B. Wegbreit, "Shape from symmetry," in *International Conference on Computer Vision*, 2005, pp. 1824–1831.

[14] O. Kroemer, H. Ben Amor, M. Ewerton, and J. Peters, "Point cloud completion using symmetries and extrusions," in *Proceedings of the International Conference on Humanoid Robots*, 2012.

[15] N. J. Mitra, L. Guibas, and M. Pauly, "Partial and approximate symmetry detection for 3d geometry," *ACM Transactions on Graphics*, vol. 25, no. 3, 2006.

[16] J. Bohg, M. Johnson-Roberson, B. León, J. Felip, X. Gratal, N. Bergström, D. Kragic, and A. Morales, "Mind the gap - robotic grasping under incomplete observation," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 686–693.

[17] M. Madry, D. Song, C. H. Ek, and D. Kragic, ""Robot bring me something to drink from": object representation for transferring task specific grasps," in *ICRA Workshop on Semantic Perception, Mapping and Exploration*, 2012.

[18] D. Song, K. Huebner, V. Kyrki, and D. Kragic, "Learning task constraints for robot grasping using graphical models," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 1579 –1585.

[19] M. Madry, D. Song, and D. Kragic, "From object categories to grasp transfer using probabilistic reasoning," in *IEEE International Conference on Robotics and Automation*, 2012, pp. 1716–1723.

[20] R. Detry, C. H. Ek, M. Madry, and D. Kragic, "Learning a dictionary of prototypical grasp-predicting parts from grasping experience," in *IEEE International Conference on Robotics and Automation*, 2013.

[21] R. Detry, C. H. Ek, M. Madry, J. Piater, and D. Kragic, "Generalizing grasps across partly similar objects," in *IEEE International Conference on Robotics and Automation*, 2012.

[22] R. Detry, C. H. Ek, M. Madry, and D. Kragic, "Compressing grasping experience into a dictionary of prototypical grasp-predicting parts," in *International Workshop on Human-Friendly Robotics*, 2012.

[23] J. Sweeney and R. A. Grupen, "A model of shared grasp affordances from demonstration," in *Humanoids*, 2007, pp. 27–35.

[24] C. Barck-Holst, M. Ralph, F. Holmar, and D. Kragic, "Learning grasping affordance using probabilistic and ontological approaches," in *International Conference on Advanced Robotics*, 2009, pp. 1–6.

[25] J. Bohg, K. Welke, B. Leon, M. Do, D. Song, W. Wohlkinger, A. Aldoma, M. Madry, M. Przybylski, T. Asfour, H. Marti, D. Kragic, A. Morales, and M. Vincze, "Task-based grasp adaptation on a humanoid robot," in *IFAC Symposium on Robot Control*, 2012, pp. 779–786.

[26] H. Dang and P. K. Allen, "Semantic grasping: Planning robotic grasps functionally suitable for an object manipulation task," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 1311–1317.

[27] M. Hanheide, C. Gretton, R. Dearden, N. Hawes, J. L. Wyatt, A. Pronobis, A. Aydemir, M. Göbelbecker, and H. Zender, "Exploiting Probabilistic Knowledge under Uncertain Sensing for Efficient Robot Behaviour," in *International Joint Conferences on Artificial Intelligence*, 2011, pp. 2442–2449.

[28] S. Hart, R. A. Grupen, and D. Jensen, "A relational representation for procedural task knowledge," in *AAAI*, 2005, pp. 1280–1285.

[29] J. Winkler, G. Bartels, L. Msenlechner, and M. Beetz, "Knowledge Enabled High-Level Task Abstraction and Execution," *Conference for Advances in Cognitive Systems*, vol. 2, no. 1, pp. 131–148, 2012.

[30] M. Tenorth and M. Beetz, "Knowrob knowledge processing for autonomous personal robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 4261–4266.

[31] Z. C. Marton, R. B. Rusu, D. Jain, U. Klank, and M. Beetz, "Probabilistic categorization of kitchen objects in table settings with a composite sensor," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 4777–4784.

[32] T. Lang and M. Toussaint, "Planning with noisy probabilistic relational rules," *Journal of Artificial Intelligence Research*, vol. 39, pp. 1–49, 2010.

[33] M. Toussaint, N. Plath, T. Lang, and N. Jetchev, "Integrated motor control, planning, grasping and high-level reasoning in a blocks world using probabilistic inference," in *IEEE International Conference on Robotics and Automation*, 2010, pp. 385–391.

[34] A. Cocora, K. Kersting, C. Plagemann, W. Burgard, and L. D. Raedt, "Learning Relational Navigation Policies," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2006, pp. 2792–2797.

[35] N. Abdo, H. Kretzschmar, and C. Stachniss, "From Low-Level Trajectory Demonstrations to Symbolic Actions for Planning," in *ICAPS Workshop on Combining Task and Motion Planning for Real-World Applications*, 2012.

[36] J. Kulick, M. Toussaint, T. Lang, and M. Lopes, "Active learning for teaching a robot grounded relational symbols," in *International Joint Conferences on Artificial Intelligence*, 2013.

[37] B. Limketkai, L. Liao, and D. Fox, "Relational Object Maps for Mobile Robots," in *International Joint Conferences on Artificial Intelligence*, 2005, pp. 1471–1476.

[38] B. Moldovan, L. Antanas, and M. Hoffmann, "Opening doors: An initial SRL approach," in *Lecture Notes in Computer Science*, vol. 7842, 2013, pp. 178–192.

[39] B. Moldovan, P. Moreno, M. van Otterlo, J. Santos-Victor, and L. De Raedt, "Learning relational affordance models for robots in multi-object manipulation tasks," in *IEEE International Conference on Robotics and Automation*, 2012, pp. 4373 –4378.

[40] M. Muja and M. Ciocarlie, "Table top segmentation package." [Online]. Available: http://www.ros.org/wiki/tabletop_object_detector

[41] M. Neumann, P. Moreno, L. Antanas, R. Garnett, and K. Kersting, "Graph kernels for object category prediction in task–dependent robot grasping," in *Eleventh Workshop on Mining and Learning with Graphs*, 2013.

[42] M. Neumann, N. Patricia, R. Garnett, and K. Kersting, "Efficient Graph Kernels by Randomization," in *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 2012, pp. 378–393.

[43] M. Datar and P. Indyk, "Locality-sensitive hashing scheme based on p-stable distributions," in *Proceedings of the 20th Annual Symposium on Computational Geometry (SCG-2004)*, 2004, pp. 253–262.

[44] J. Vennekens, M. Denecker, and M. Bruynooghe, "CP-logic: A language of causal probabilistic events and its relation to logic programming," *Theory and Practice of Logic Programming*, vol. 9, no. 3, pp. 245–308, 2009.

[45] W. Meert, J. Struyf, and H. Blockeel, "Learning ground CP-logic theories by leveraging bayesian network learning techniques," *Fundam. Inform.*, vol. 89, no. 1, pp. 131–160, 2008.

[46] D. Fierens, G. Van den Broeck, I. Thon, B. Gutmann, and L. De Raedt, "Inference in probabilistic logic programs using weighted CNF's," in *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence*, F. Gagliardi Cozman and A. Pfeffer, Eds., July 2011, pp. 211–220.

[47] D. Fischinger, M. Vincze, and Y. Jiang, "Learning grasps for unknown objects in cluttered scenes," in *IEEE International Conference on Robotics and Automation*, 2013.

[48] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, pp. 273–297, 1995.

[49] J. C. Platt, "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," in *Advances in Large Margin Classifiers*. MIT Press, 1999, pp. 61–74.

[50] Y. Jiang, S. Moseson, and A. Saxena, "Efficient grasping from rgbd images: Learning using a new rectangle representation," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 3304–3311.

[51] H. Baltzakis, "Orca simulator." [Online]. Available: http://www.ics.forth.gr/cvrl/_software/orca_setup.exe

[52] ASUS, "Asus xtion pro." [Online]. Available: http://www.asus.com/Multimedia/Motion_Sensor/Xtion_PRO/

[53] W. robotics, "Universal gripper WSG 50." [Online]. Available: http://www.weiss-robotics.de/gripper-systems/gripper-modules/universal-gripper-wsg-50.html

[54] K. robotics, "Kuka lightweight robot (LWR)." [Online]. Available: http://www.kuka-robotics.com/en/products/addons/lwr/

[55] J. Jerez and A. Suero, "Newton Game Dynamics. Open-source Physics Engine." [Online]. Available: http://www.newtondynamics.com

[56] G. Sánchez and J.-C. Latombe, "A Single-Query Bi-Directional Probabilistic Roadmap Planner with Lazy Collision Checking," in *Robotics Research*, ser. Springer Tracts in Advanced Robotics. Springer Berlin Heidelberg, 2003, vol. 6, pp. 403–417.

[57] I. A. Şucan, M. Moll, and L. E. Kavraki, "The Open Motion Planning Library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.

[58] X. Zhu, "Persistent homology: An introduction and a new text representation for natural language processing," in *International Joint Conferences on Artificial Intelligence*, 2013.

[59] B. Moldovan, P. Moreno, and M. van Otterlo, "On the use of probabilistic relational affordance models for sequential manipulation tasks in robotics," in *IEEE International Conference on Robotics and Automation*, 2013.

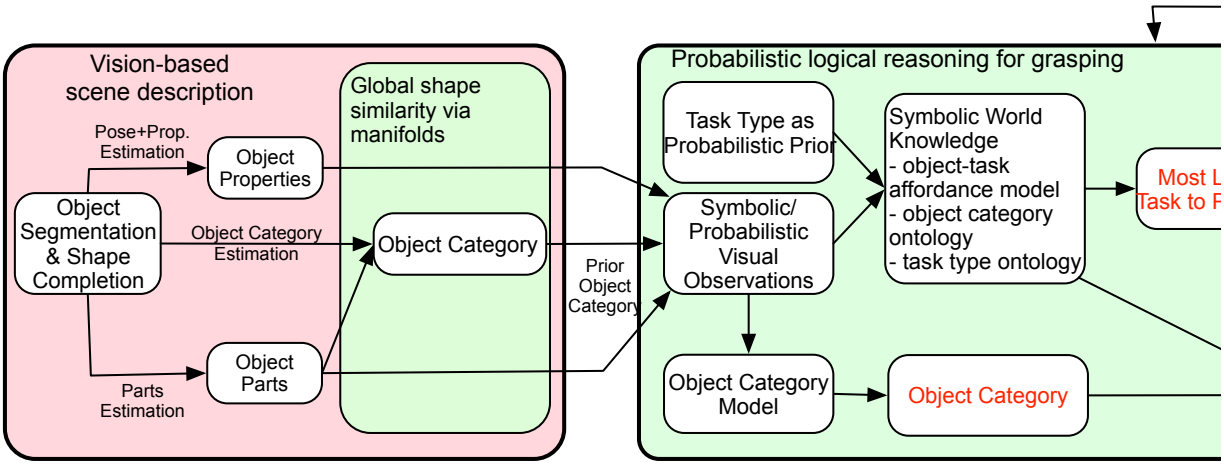

Fig. 2: The task-dependent grasping pipeline.